

TITLE OF THE INVENTION

IMAGE ENCODING METHOD AND APPARATUS

5

FIELD OF THE INVENTION

This invention relates to an image encoding method and apparatus for encoding an input image by applying quantization processing that differs for each region of the image.

BACKGROUND OF THE INVENTION

Recent advances in digital signal processing technology have made it possible to efficiently encode large quantities of digital information such as moving and still pictures and video and to record the encoded information on a small-size magnetic medium or to transmit it to a communication medium.

A technique using the discrete wavelet transform is known as a highly efficient method of encoding an image. In accordance with this technique, the discrete wavelet transform is applied to an input image signal to be encoded. In the discrete wavelet transform, two-dimensional discrete wavelet transform processing is applied to an input image signal, and then a sequence of

coefficients obtained by the discrete wavelet transform is quantized.

In such quantization, a region of an image to be encoded to an image quality higher than that of a peripheral portion of an image containing the image region is designated by a user. The coefficients that belong to the designated region are then evaluated, these coefficients are quantized upon raising the precision of quantization a prescribed amount, and encoding is carried out in such a manner that the designated image region can be decoded to an image quality higher than that of the periphery.

With this conventional technique, however, the designation of the image region desired to be encoded to a high image quality is an explicit designation made by the user. The operation demanded of the user is therefore a complicated one.

Further, if it is so arranged that the image region to thus be encoded to a high image quality is determined by automatically discriminating the patterns or colors of this image, a limitation is imposed on the colors or shapes of objects to be encoded to the high image quality and it will not be possible to obtain an object that can be used universally. For example, in a case where video shot by a home digital video camera or the like is to be processed, satisfactory results are not

obtained.

Further, the specification of Japanese Patent Application Laid-Open No. 10-145606 describes a region discrimination method as a technique through which a
5 wavelet transform is applied to an input image and a region of interest in the image is extracted using subband signals that are obtained. According to the invention described in this publication, separation of an image region is implemented depending upon whether a
10 wavelet coefficient obtained by applying a Harr wavelet transform to an image signal, i.e., the absolute value of the high-frequency component of the subband signals, is greater than a predetermined threshold value.

With this example of the prior art, however, the
15 purpose is to separate a region having a strong edge from a region having a weak edge by referring to the absolute values of the wavelet coefficients (i.e., of the subband signals). The segmentation of a higher-order multilevel region or the extraction of a region of
20 interest, namely the extraction of a subject of interest from an image region, cannot be carried out.

SUMMARY OF THE INVENTION

25 An object of the present invention is to provide an image encoding method and apparatus through which

diverse image regions can be designated and encoded efficiently without placing a burden upon the user.

Another object of the present invention is to provide an image encoding method and apparatus through
5 which an image region to be encoded to a higher level can be selected and encoded automatically in accordance with the characteristics of the image to be encoded.

A further object of the present invention is to provide an image encoding method and apparatus through
10 which regions of interest are extracted from image data automatically and encoding processing that differs for each extracted region can be executed.

In order to attain the above described objects, an
15 image encoding apparatus of the present invention comprises: image input means for inputting an image signal; band dividing means for dividing the image signal input by said image input means into different spatial frequency bands; region-of-interest extraction
20 means for extracting a region of interest by obtaining a distribution of motion vectors in the image signal based upon values of spatial frequency components of the image signal obtained by the band dividing means; quantization means for applying quantization processing to the region
25 of interest extracted by the region-of-interest extraction means and different quantization processing

to other regions, and outputting a quantized image signal; and image encoding means for encoding the quantized image signal quantized by the quantization means.

5 In order to attain the above described objects, an image encoding apparatus of the present invention comprises: transformation means for applying a discrete wavelet transform to an image signal; motion detection means for detecting motion of an image based upon the
10 image signal; region designation means for designating a region of the image signal based upon information indicating motion of the image detected by the motion detection means; quantization means for quantizing a discrete wavelet transformed output from the
15 transformation means in accordance with the region designated by the region designation means and outputting a quantized image signal; and encoding means for encoding the quantized image signal quantized by the quantization means.

20 Other features and advantages of the present invention will be apparent from the following description taken in conjunction with the accompanying drawings, in which like reference characters designate the same or similar parts throughout the figures
25 thereof.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and constitute a part of the specification,
5 illustrate embodiments of the invention and, together with the description, serve to explain the principle of the invention.

Fig. 1 is a block diagram illustrating the construction of an image encoding apparatus according to
10 a first embodiment of the present invention;

Figs. 2A and 2B are diagram useful in describing a wavelet transform in a discrete wavelet transformation unit according to the first embodiment;

Fig. 3 is a block diagram illustrating the construction of an ROI extraction unit according to the
15 first embodiment;

Fig. 4 is a block diagram illustrating the construction of a motion vector detector according to the first embodiment;

20 Figs. 5A, 5B and 5C are diagrams useful in describing an ROI mask and a quantization method according to the first embodiment;

Fig. 6 is a diagram useful in describing entropy encoding;

Fig. 7 is a block diagram illustrating the construction of an ROI extraction unit according to a second embodiment of the present invention;

Fig. 8 is a diagram useful in describing the notation of an equation for calculating degree of left-right symmetry according to the second embodiment;

Fig. 9 is a block diagram illustrating the construction of an image encoding apparatus according to a fourth embodiment of the present invention;

Fig. 10 is a block diagram illustrating the construction of a motion vector detector according to the fourth embodiment;

Fig. 11 is a block diagram illustrating the construction of a motion vector detector according to a fifth embodiment;

Fig. 12 is a block diagram illustrating the construction of a motion vector detector according to a sixth embodiment; and

Fig. 13 is a block diagram illustrating the construction of an region designation unit according to a seventh embodiment.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Preferred embodiments of the present invention will now be described in detail with reference to the

accompanying drawings.

[First Embodiment]

Fig. 1 is a block diagram illustrating the construction of an image encoding apparatus according to a first embodiment of the present invention.

As shown in Fig. 1, the apparatus includes an image input unit 1 for inputting image data. By way of example, the image input unit 1 is equipped with a scanner for reading a document image, with an imaging device such as a digital camera, or with an interface for interfacing a communication line. The input image is applied to a discrete wavelet transformation unit 2, which applies a two-dimensional discrete wavelet transform to the input image. An ROI (Region of Interest) extraction unit 3 extracts an ROI from the image that has entered from the image input unit 1. A quantizer 4 quantizes coefficients obtained by the two-dimensional discrete wavelet transform. An encoder 5 encodes the image signal that has been quantized by the quantizer 4, and a code output unit 6 outputs the code obtained by the encoder 5.

The apparatus according to the first embodiment is not limited to a special-purpose apparatus of the kind shown in Fig. 1 and is applicable also to a case where a program which implements these functions is loaded in, e.g., a general-purpose personal computer or work

station and the computer or work station is made to operate in accordance with the program.

The operation of the apparatus will now be described with reference to Fig. 1.

5 First, an image signal constituting an image to be encoded is input to the image input unit 1 by raster scanning. The input thus entered is input to the discrete wavelet transformation unit 2. In the description that follows, it will be assumed that the
10 image signal that has entered from the image input unit 1 is a monochrome multilevel image. However, if an image signal having a plurality of color components, such as a color image, is input and encoded, it will suffice to compress the RGB color components or the
15 luminance and chromaticity components as well as the monochrome components.

The discrete wavelet transformation unit 2 subjects the input image signal to two-dimensional discrete wavelet transform processing, calculates the transform
20 coefficients and outputs these coefficients. The first embodiment assumes application of the Haar wavelet transform, which best lends itself to hardware implementation. A low-pass filter (referred to as an "LPF" below) employed in the Haar wavelet transform
25 averages mutually adjacent pixels, and a high-pass filter (referred to as an "HPF" below) calculates the

difference between the mutually adjacent pixels.

The procedure of two-dimensional discrete wavelet transform processing will be described with reference to Figs. 2A and 2B.

5 Fig. 2A is a diagram useful in describing horizontal- and vertical-direction transform processing applied to an input image signal. Filtering by an LPF and an HPF is performed first in the horizontal direction. A sequence of low-pass coefficients and a
10 sequence of high-pass coefficients thus obtained are each downsampled, to half the rate, in the horizontal direction by downsamplers 201. Next, filtering similar to that in the horizontal direction is applied in the vertical direction and then downsampling to half the
15 rate is applied by downsamplers 202 in the vertical direction. By repeatedly executing the same processing to signals of the lowest frequency band, eventually a series of data sequences (LL, LH2, HL2, HH2, LH1, HL1, HH1) belonging to seven different frequency bands are
20 output.

Fig. 2B illustrates the manner in which an input multilevel image signal is divided into different frequency bands as a result of the transform processing shown in Fig. 2A.

25 As shown in Fig. 2B, the frequency bands are labeled HH1, HL1, LH1, ..., LL. In the description that

follows, a single transformation in the horizontal and vertical directions shall be considered to be one level of decomposition, and the frequency bands HH1, HL1, LH1, ..., LL shall be referred to as "subbands". The principle of image compression based upon this wavelet transform is reported in detail in M. Antonini, M. Barlaud, P. Mathieu and I. Daubechies, "Image Coding Using Wavelet Transform", IEEE Transactions on Image Processing, Vol. 1, No. 2, April 1992.

Fig. 3 is a block diagram useful in describing the construction of the ROI extraction unit 3 according to the first embodiment.

As shown in Fig. 3, the ROI extraction unit 3 includes a motion vector detector (MVD) 10 and a region segmentation unit 11. Subbands obtained by dividing the image signal into the frequency bands using the discrete wavelet transformation unit 2 enter the motion vector detector 10.

Motion-vector estimation is performed based upon the well-known gradient method (also referred to as the temporal-spatial gradient method or temporal-spatial differentiation method, etc.) For a description of the principle of the gradient method, see USP 3,890,462; J.O. Limb and J.A. Murphy, "Measuring the Speed of Moving Objects from Television Signals", IEEE Transactions on Communications, Vol. COM23, pp. 474 -

478, April 1975; and J.O. Limb and J.A. Murphy,
"Estimating the Velocity of Moving Images in Television
Signals", Computer Graphics and Image Processing, 4, pp.
311 - 327, 1975. Equations for estimating motion

5 vectors based upon the gradient method are as follows:

$$\alpha = -\sum B\{\Delta t(i) \cdot \text{sign}(\Delta x(i))\} / \sum B|\Delta x(i)| \quad \dots (1)$$

$$\beta = -\sum B\{\Delta t(i) \cdot \text{sign}(\Delta y(i))\} / \sum B|\Delta y(i)| \quad \dots (2)$$

where α and β represent the results of estimating, in
the horizontal and vertical directions, respectively, a
10 motion vector V at a pixel of interest, $\Delta t(i)$ represents
the amount of change with time of a pixel value of an i -
th pixel neighboring the pixel of interest, $\Delta x(i)$
represents a horizontal spatial gradient at the i -th
pixel neighboring the pixel of interest, and $\Delta y(i)$
15 represents a vertical spatial gradient at the i -th pixel
neighboring the pixel of interest. Further, $\text{sign}(x)$
represents an operator for extracting the sign bit of an
input signal x , and $|x|$ represents an operator for
outputting the absolute value of the input signal x . In
20 addition, $\sum B$ represents the sum total within a block B
comprising a plurality of pixels centered on the pixel
of interest. The motion vector $V(\alpha, \beta)$ at the pixel of
interest is estimated using the temporal change $\Delta t(i)$,
horizontal spatial gradient $\Delta x(i)$ and vertical spatial
25 gradient $\Delta y(i)$ of pixel values of all pixels i that
belong to the block B . The size of the block B in this

case is usually 3×3 to 15×15 pixels.

Fig. 4 is a block diagram showing the details of the motion vector detector 10.

As shown in Fig.4, the motion vector detector 10 includes an input unit 20 for inputting a subband LL; an input unit 21 for inputting a subband HL2 or LH2; an image memory 22 such as a frame memory; an adder 23 for performing addition or subtraction; a sign output unit (sign) 24 for extracting the sign bit of input data; a multiplier 25; an absolute-value output unit (ABS) 26 for outputting the absolute value of input data; accumulators (ΣB) 27, 28 for performing cumulative addition; a divider 29 for executing division; and an output unit 30 for outputting the estimated value of a motion vector.

The subband LL that has entered from the input unit 20 is subtracted from the subband LL of the preceding frame, which has arrived via the image memory 22, by the adder 23, whereby temporal change Δt of the pixel value is calculated. Meanwhile, the subband HL2 or LH2 enters directly from the input unit 21, taking note of the fact that the horizontal and vertical spatial gradients Δx , Δy of the image have already been operated on as subbands HL2, LH2, respectively. The sign of each pixel of the subband HL2 or LH2 is output from the sign output unit 24 and applied to the multiplier 25. The latter

multiplies the temporal change Δt of the pixel value by the sign of the spatial gradient that enters from the input unit 21. The absolute-value output circuit 26 calculates the absolute value of the pixel value of each pixel in the entered subband HL2 or LH2. From the block comprising the plurality of neighboring pixels centered on the pixel of interest, the accumulators 27, 28 cumulatively add the values (the outputs of the multiplier 25) obtained by multiplying the temporal change Δt by the sign of the spatial gradient, and the absolute values (the outputs of the absolute-value output circuit 26) of the spatial gradient Δx or Δy , respectively. More specifically, the accumulator 27 calculates the numerators of Equations (1), (2) and the accumulator 28 calculates the denominators of Equations (1), (2). Finally, the divider 29 performs the division in accordance with Equations (1), (2) and the output unit 30 outputs the horizontal component α or vertical component β of the motion vector. In accordance with the procedure described above, a minute distribution of motion vectors can be obtained over the entire area of the input image.

Next, the region segmentation unit 11 subjects the image to region segmentation by referring to the distribution of motion vectors detected by the motion vector detector 10. Within the image to be encoded, a

region (ROI) that is to be decoded at a quality higher than that at the image periphery is decided and mask information indicating which coefficients belong to the designated region is generated when the image of

5 interest is subjected to the discrete wavelet transform. It should be noted that the determination of the ROI can be performed by referring to the picture-taking mode of the camera. For example, if the camera is in a tracking photographic mode, the subject (the ROI) is

10 substantially stationary at the center of the picture and the background travels in conformity with the motion of the subject. If the camera is set to a mode for photography using a tripod, the subject (the ROI) will move freely within the picture and the background will

15 be substantially stationary. Accordingly, which region in the input image is the present ROI can be determined from the mode of photography.

Fig. 5A is a conceptual view for describing an example of a case where mask information for extracting only the ROI or for excluding only the ROI is generated.

If a star-shaped ROI exists in an image, as indicated on the left side of Fig. 5A, the region segmentation unit 11 extracts the ROI based upon the motion-vector distribution information and calculates the portion of each subband occupied by this ROI. The region indicated by the mask information is a domain,

which includes the transform coefficients of the
periphery, necessary when decoding the image signal on
the boundary of the ROI. An example of the mask
information thus calculated is shown on the right side
5 of Fig. 5A. In this example, mask information for when
a two-level two-dimensional discrete wavelet transform
is applied to the image on the left side of Fig. 5A is
calculated in the manner shown on the right side of Fig.
5A. In Fig. 5A, the star-shaped portion is the ROI, the
10 bits constituting the mask information within this
region are "1"s and the bits of the other region
information are "0"s. The entirety of this mask
information is identical with the constitution of the
transform coefficients obtained by the two-dimensional
15 discrete wavelet transform. By scanning the bits within
the mask information, therefore, it is possible to
identify whether the coefficients at the corresponding
positions fall within the designated region.

The mask information thus produced is applied to
20 the quantizer 4. Furthermore, the region segmentation
unit 11 receives an input of a parameter, which
specifies the image quality of the designated ROI, from
an input unit (e.g., a keyboard or a pointing device
such as a mouse), which is not shown. The parameter may
25 be a numerical value expressing a compression rate
assigned to the designated region, or a numerical value

representing the image quality of this region. On the basis of this parameter, the region segmentation unit 11 calculates a bit-shift quantity W for the coefficients in the ROI and outputs this to the quantizer 4 together with the mask information.

The quantizer 4 quantizes the transform coefficients from the discrete wavelet transformation unit 2 by a predetermined quantization step Δ and outputs indices corresponding to the quantized values.

10 Quantization is carried out in accordance with the following equations:

$$q = \text{sign}(c) \cdot \text{floor}(|c|/\Delta) \quad \dots (3)$$

$$\text{sign}(c) = -1; \quad c < 0 \quad \dots (5)$$

Next, the quantizer 4 changes the quantization indices in accordance with the following equations based upon the mask information and shift quantity W that has entered from the ROI extraction unit 3:

$$q' = q \quad ; \quad m = 0 \quad \dots (7)$$

where m represents the value of mask information at the position of the quantization index. By virtue of the processing described above, only a quantization index
 5 that belongs to the designated spatial region is shifted up by W bits in the ROI extraction unit 3.

Figs. 5B and 5C are diagrams useful in describing a change in quantization index by such shift-up.

In Fig. 5B, three quantization indices exist in
 10 three subbands. If the value of mask information of a quantization index that has been subjected to screening is "1" and the number W of shifts is "2", then the quantization indices after the shift will be as shown in Fig. 5C. The quantization indices that have been
 15 changed in this manner are output to the encoder 5.

In this embodiment, entropy encoding is used as the encoding method in the encoder 5. Entropy encoding will be described below.

The encoder 5 decomposes entered quantization
 20 indices into bit planes, applies binary arithmetic encoding on a per-bit-plane basis and outputs a code sequence.

Fig. 6 is a diagram useful in describing the operation of the encoder 5. In this example, three non-
 25 zero quantization indices exist in a region within a subband having a size of 4×4 , and the values of these

indices are "+13", "-6" and "+3". The encoder 5 obtains a maximum value M by scanning this region and, in accordance with the following equation, calculates a number S of bits necessary to express the maximum

5 quantization index:

$$S = \text{ceil}(\log_2(|M|)) \quad \dots (8)$$

where $\text{ceil}(x)$ is a function for outputting the smallest integral value that is greater than x.

In Fig. 6, the maximum coefficient value is "13"
10 and therefore the value of S is "4" and the 16 quantization indices in the sequence are processed in units of the four bit planes, as shown in Fig. 6.

First, the entropy encoder 5 applies binary arithmetic encoding to each bit of the most significant
15 bit plane (represented by MSB in Fig. 6) and outputs the encoded bits as a bit stream. Next, the bit plane is lowered by one level and the process is repeated. This processing is repeated until the bit plane of interest reaches the least significant bit plane (represented by
20 LSB), with each bit of the bit planes being encoded and output to the code output unit 6. If an initial non-zero bit is detected in the scanning of the bit planes, then the code of this quantization index undergoes entropy encoding immediately thereafter. The encoded
25 image signal is finally output from the code output unit 6.

Thus, in accordance with the first embodiment as described above, the following effects are obtained:

(1) Using a subband signal obtained by applying a wavelet transform to an input image makes it possible to
5 estimate a motion vector within this input image.
Accordingly, the wavelet transform can be used not only for encoding an image but also for ROI extraction without requiring a large-scale modification or increase in hardware.

10 (2) As a result, an image signal can be compressed more efficiently as by using a certain quantization step width for an extracted ROI and a different quantization step width for a non-ROI.

(3) Since ROI extraction is performed using a
15 subband signal of a reduced sampling rate as the target, processing is quick.

In the first embodiment, the Haar wavelet transform is applied to an image signal. However, similar results can be obtained with any wavelet transform in which the
20 high-pass filter (HPF) reflects the spatial gradient of the image.

Further, in the wavelet transform, the same filtering processing is repeatedly applied to the subband of the minimum frequency of the input image,
25 whereby the input image is converted to multiple resolutions in a pyramid structure. Accordingly,

detection of a motion vector in the first embodiment means not only extraction independently from a subband signal in a specific region, as described above. For example, it is possible to estimate a motion vector
5 rapidly at a coarse resolution and, by referring to this low-resolution motion vector, to precisely estimate motion vectors in subband signals having a gradually higher resolution.

[Second Embodiment]

10 In the first embodiment described above, motion vectors within an image are detected minutely using subband signals obtained by application of the Haar wavelet transform, and an ROI is extracted based upon the distribution of these motion vectors. In the second
15 embodiment, an ROI having left-right symmetry is extracted within an image having substantial left-right symmetry, such as an image of the human face, using subband signals obtained by application of the Haar wavelet transform in a manner similar to that of the
20 first embodiment.

Fig. 7 is a block diagram illustrating the construction of the ROI extraction unit 3 according to the second embodiment for extracting an ROI having left-right symmetry. The ROI extraction unit 3 includes an
25 arithmetic unit 40 for calculating degree of left-right symmetry and a region segmentation unit 41.

As shown in Fig. 7, in subbands obtained from the preceding two-dimensional discrete wavelet transformation unit 2, HL2 having horizontal spatial-gradient information and LH2 having vertical spatial-gradient information are input to the arithmetic unit 40 for calculating degree of left-right symmetry. A method in which the similarity of brightness or color on the left and right sides of the ROI is evaluated as a yardstick of left-right symmetry is a direct approach. However, since the luminance and color of an image are readily influenced by illumination, this evaluation method does not necessarily provide stable results. Accordingly, in the second embodiment, information concerning the orientation of the spatial gradient of an image is utilized as an indicator that is not readily susceptible to the lighting conditions. (For example, see the specification of Japanese Patent Application Laid-Open No. 10-162118 or Toshiaki Kondo and Hong Yan, "Automatic Human Face Detection and Recognition under Non-uniform Illumination", Pattern Recognition, Vol. 32, No. 10, pp. 1707 - 1718, October 1999.)

Extraction of a region having left-right symmetry utilizing information relating to the orientation of the spatial gradient of an image will be considered. To accomplish this, a region of interest having a regular shape, such as a rectangular or elliptical block, is set

within an input image, the region of interest is shifted in position incrementally within the input image and, each time a shift is made, the degree of left-right symmetry is calculated. Since the region of interest
5 has left-right symmetry, the spatial gradients of pixels at corresponding positions on the left and right sides of a perpendicular axis which crosses this region must satisfy the following two conditions:

(i) the orientations of the gradients in the
10 horizontal direction are opposite each other; and

(ii) the orientations of the gradients in the vertical direction are identical.

Orientation θ of a spatial gradient is expressed by Equation (9) below as the ratio of gradient Δx in the
15 horizontal direction to gradient Δy in the vertical direction at each pixel.

$$\theta(x,y) = \tan^{-1}(\Delta y/\Delta x) \quad \dots (9)$$

where " \tan^{-1} " represents arc-tangent (arctan). Taking note of the fact that the horizontal and vertical
20 spatial gradients Δx , Δy of the image have already been calculated as subbands HL2, LH2, respectively, Equation (9) can be rewritten as Equation (10) below.

$$\theta(x,y) = \tan^{-1}(LH2/HL2) \quad \dots (10)$$

According to the second embodiment, therefore, the
25 orientation of a spatial gradient is found by implementing Equation (10) making direct use of the

outputs LH2, HL2 of the discrete wavelet transformation
 unit 2. Though Equation (10) may be evaluated as is
 using real numbers, the operation can be executed more
 simply and at higher speed if a look-up table is
 5 utilized. Next, degree γ of symmetry at each position
 (x,y) of the image is found. According to the second
 embodiment, $\gamma(x,y)$ is defined as indicated by Equation
 (11) below using $\theta(x,y)$.

$$r(x,y) = \sum_{j=y-v/2}^{y+v/2} \sum_{i=x-h/2}^x [|\cos(\theta(i,j)) + \cos(\theta(2x-i,j))|^2 + |\sin(\theta(i,j)) - \sin(\theta(2x-i,j))|^2]$$

10 ... (11)

Equation (11) will be described with reference to
 Fig. 8, which illustrates a rectangular block to be
 operated on in order to calculate the degree γ of
 symmetry.

15 In Fig. 8, **h** represents the size of the block in
 the horizontal direction, **v** the size of the block in the
 vertical direction, and (x,y) the position of the center
c of the block. Calculation of the degree γ of symmetry
 is performed while raster-scanning a pixel of interest
 20 from a pixel **s** $(x-h/2,y-v/2)$ to a pixel **e** $(x,y+v/2)$
 within the left half of the block in the manner
 illustrated. For example, when the pixel of interest is
m, the degree γ of symmetry is calculated using spatial-
 gradient information at the pixel **m** and the amount of
 25 the spatial gradient at a pixel **n**, which is located at a

position symmetrical with respect to the pixel m . If the degree of left-right symmetry is high, the first and second terms of Equation (11) both become small values owing to cancellation within the terms. Though the
5 absolute values are squared in Equation (11), the degree γ of symmetry may simply be the sum of the absolute values. Though the size of the area to be operated on approximately agrees with the size of a face to be detected, a block having a size that differs in several
10 stages may be used if the size of the face is unknown.

The region segmentation unit 41 extracts only a region having high degree of left-right symmetry based upon the results from the arithmetic unit 40 for calculating the degree of left-right symmetry and then
15 decides the region that corresponds to the human face from the extracted region. In order to decide the position of the human face from the region of high left-right symmetry, any method whose purpose is to extract a human face may be employed, such as making combined use
20 of template matching using the template of a human face, a search for a flesh-tone region or elliptical region and motion-vector information obtained in accordance with the first embodiment.

Since extraction of an ROI in the second embodiment
25 is such that a specific region having left-right symmetry is extracted in an efficient manner, the method

of extraction is highly effective when photographing a human being. Accordingly, in a case where the image input unit 1 is a camera or the like, the function of the second embodiment is turned on when a mode for photographing a human being or a portrait mode is selected in association with the picture-taking mode. When another type of photographic mode is in effect, a default is set so as to turn the function of the second embodiment off.

10 In accordance with the second embodiment, as described above, the following advantages are obtained:

(1) A region in an input image having left-right symmetry can be extracted in a simple manner using subband signals obtained by applying a wavelet transform to the input image. Accordingly, the wavelet transform can be exploited not only for encoding an image but also for ROI extraction without requiring a large-scale modification or increase in hardware.

(2) As a result, an image signal can be compressed more efficiently as by using a certain quantization step width for an extracted ROI and a different quantization step width for a non-ROI.

(3) Since ROI extraction is performed using a subband signal of a reduced sampling rate as the target, processing is quick.

(4) The arithmetic unit 40 for calculating degree

of left-right symmetry utilizes only the gradient direction of an image. As a result, there is little influence from changes in lighting conditions and a region of left-right symmetry can be detected in stable
5 fashion. Further, since a specific region having left-right symmetry is found, regions that are candidates for a human face can be narrowed down efficiently regardless of whether the background is complicated or simple.

(5) By determining the position of a face through
10 application of pattern matching solely to portions having a high degree of left-right symmetry, a human face can be detected highly precisely and at high speed while greatly reducing the amount of pattern-matching processing.

15 In the second embodiment, the Haar wavelet transform is utilized. However, similar results can be obtained also if the HPF used in the Haar wavelet transform is of the quadratic differential type, such as a Laplacian filter.

20 [Third Embodiment]

In the first embodiment described above, motion vectors within an image are detected minutely utilizing subband signals obtained by application of the Haar wavelet transform, and an ROI is extracted based upon
25 the distribution of these motion vectors. In the second embodiment, an ROI having left-right symmetry is

extracted in similar fashion using subband signals
obtained by application of the Haar wavelet transform.
Next, in the third embodiment, segmentation of an input
image into regions is carried out utilizing subband
5 signals obtained by application of a wavelet transform
and a region of interest in the input image is
extracted.

An image having a low resolution often is used to
perform region segmentation efficiently. The reason for
10 this is that a region within an image contains more
global image information, unlike edge information, for
example. Accordingly, first the input image is
subjected to region segmentation coarsely using the
subband signal LL of the lowest frequency band.
15 According to the third embodiment, the quad-tree
segmentation method is utilized for region segmentation.
However, the present invention is not limited to the
quad-tree segmentation method. Any other method such as
the clustering method or histogram-base technique may be
20 used.

Quad-tree segmentation segments an image into
regions through the following steps:

(1) Processing is started using the input image,
namely the subband signal LL in this embodiment, as one
25 image region.

(2) Homogeneity, e.g., a variance value, within

the region is calculated. If the variance value exceeds a fixed value, the region is judged to be non-homogeneous and is split into four equal regions.

(3) If mutually adjacent split regions satisfy the
5 condition for homogeneity, i.e., exhibit variance values that are equal to or less than the fixed value, these regions are merged.

(4) The above-described steps are repeated until splitting and merging no longer occur.

10 Thus, an input image can be segmented into regions coarsely on a block-by-block basis. If an input image consists of a comparatively simple pattern, calculation can be simplified using the difference between maximum and minimum values in segmented regions as a criterion
15 for evaluating homogeneity.

Next, edge distribution $\text{edge}(i,j)$ is found using subband signals LH2, HL2.

Edge strength can expressed by the following equation:

20
$$\text{edge}(i,j) = |\text{LH2}(i,j)| + |\text{HL2}(i,j)| \dots (12)$$

As a result of the quad-tree segmentation described above, the boundary of this segmented region generally consists of connected minute blocks. Accordingly, in a region in which these minute blocks are connected,
25 pixels of a strong edge defined by Equation (12) are traced and the resulting path is decided upon as being

the true boundary line.

According to the third embodiment, as described above, a region of interest cannot be specified even though an image can be segmented. However, the

5 following effects can be obtained by combining the third embodiment with the first and/or second embodiments or by adding on initial information concerning a region of interest:

(1) In a case where the input image is a sequence
10 of moving pictures, a region of interest can be specified from the motion-vector distribution, which was described in the first embodiment, and the picture-taking mode, and the contour of this region of interest can be finalized as the true contour through the
15 procedure set forth above.

(2) In a case where a human face is to be extracted, a boundary line defining an elliptical shape can be finalized as the contour through the above-described procedure by focusing upon the axes of
20 symmetry of a symmetric region as described in the second embodiment.

(3) If a region of interest is designated at the start in a case where the input image is a sequence of moving pictures, then the contour of this region of
25 interest can be traced by repeating the above-described procedure.

[Fourth Embodiment]

Fig. 9 is a block diagram illustrating an example of the construction of an image encoding apparatus according to a fourth embodiment of the present invention.

As shown in Fig. 9, the apparatus includes an image input unit 101 for inputting image data. By way of example, the image input unit 101 is equipped with a scanner for reading a document image, with an imaging device such as a digital camera, or with an interface for interfacing a communication line. The input image is applied to a discrete wavelet transformation unit 102, which applies a two-dimensional discrete wavelet transform to the input image. A quantizer 103 quantizes a sequence of transform coefficients obtained by the two-dimensional discrete wavelet transformation unit 102, and an entropy encoder 104 applies entropy encoding to the image signal quantized by the quantizer 103. A code output unit 105 outputs the code obtained by the encoder 104. A motion detector 107 detects the motion of an object in the image that has entered from the image input unit 101. On the basis of the motion of the object in the image detected by the motion detector 107, a region designation unit 106 determines a region to be subjected to a particularly high degree of encoding, sends the result of determination to the quantizer 103

and instructs the quantizer 103 to perform quantization.
The components 101 to 106 in Fig. 9 correspond to the
components 1 to 6, respectively, in Fig. 1. Though the
motion vector detector 10 shown in Fig. 3 detects the
5 motion of an image based upon subbands after application
of the wavelet transform, the motion detector 107 in
this embodiment detects motion based upon the original
image signal. Further, the apparatus according to the
fourth embodiment is not limited to a special-purpose
10 apparatus of the kind shown in Fig. 9 and is applicable
also to a case where a program which implements these
functions is loaded in, e.g., a general-purpose personal
computer or work station and the computer or work
station is made to operate in accordance with the
15 program.

The operation of the apparatus will now be
described with reference to Fig. 9.

First, an image signal constituting an image to be
encoded is input to the image input unit 101 by raster
20 scanning. The input thus entered is input to the
discrete wavelet transformation unit 102 and to the
motion detector 107. In the description that follows,
it will be assumed that the image signal that has
entered from the image input unit 101 is a monochrome
25 multilevel image. However, if an image signal having a
plurality of color components, such as a color image, is

input and encoded, it will suffice to compress the RGB color components or the luminance and chromaticity components as the monochrome components.

The discrete wavelet transformation unit 102
5 subjects the input image signal to two-dimensional discrete wavelet transform processing and applies the sequence of coefficients resulting from the transformation to the quantizer 103. As is well known, a two-dimensional discrete wavelet transform can be
10 expressed by successively applying a one-dimensional discrete wavelet transform successively in the horizontal and vertical directions of the image. A one-dimensional discrete wavelet transform divides the input signal into low- and high-frequency components by
15 prescribed low- and high-pass filters and downsamples each of these components to half the number of samples.

On the basis of the image signal supplied by the image input unit 101, the motion detector 107 detects a region of motion within the image and supplies the
20 region designation unit 106 with a detection signal 110 indicative of the result of detection. When the detection signal 110 enters, the region designation unit 106 outputs region information 111, which is for instructing the quantizer 103 to execute highly
25 efficient encoding.

Fig. 10 is a block diagram illustrating an example

of the construction of the motion detector 107 according to the fourth embodiment. This arrangement is applied to a case where the input image signal is an interlaced image signal typified by a television signal.

5 As shown in Fig. 10, the motion detector 107 includes line delay circuits 201, 202 and a comparator 203. The image signal from the image input unit 101 is supplied to the comparator 203 along a total of three paths, namely a path $P(x, y+1)$ leading directly to the
10 comparator 203, a path $P(x, y)$ via the line delay circuit 201 and a path $P(x, y-1)$ via both line delay circuits 201 and 202. The line delay circuits 201 and 202 are each single-pixel delay circuits corresponding to one horizontal line of the image signal. Accordingly, sets
15 of three pixels arrayed vertically are supplied to the comparator 203 sequentially. The comparator 203 compares the average value of the upper and lower pixels of the three vertically arrayed pixels of the set with the value of middle pixel and determines whether the
20 difference between the compared values exceeds a predetermined quantity. More specifically, the comparator 203 detects motion between fields in the interlaced image signal and supplies the result to the region designation unit 106.

25 According to the fourth embodiment, the detection signal 110 is output as a high-level signal if the

following relation holds:

$$\text{abs}\{(x,y+1) + p(x,y-1))/2 - P(x,y)\} > K \quad \dots (13)$$

where K represents a predetermined value. It should be noted that $\text{abs}\{(x,y+1) + p(x,y-1))/2 - P(x,y)\}$ in

5 Equation (13) indicates the absolute value of the difference between the value of $P(x,y)$ and the average of the values of pixel $P(x,y+1)$ and pixel $P(x,y-1)$.

In accordance with the fourth embodiment as described above, motion of an image can be detected
10 automatically based upon the difference between vertically arrayed pixel values contained in the image, thereby making it possible to select an image region to undergo highly efficient encoding.

[Fifth Embodiment]

15 A fifth embodiment of the invention in which the motion detector 107 has a different construction will now be described.

Fig. 11 is a block diagram illustrating an example of the construction of a motion vector detector 107a
20 according to a fifth embodiment. This arrangement is used in a case where the input image signal is a progressive image signal typified by an image signal processed by a personal computer or the like.

As shown in Fig. 11, the motion detector 107a
25 includes a frame delay circuit 301 for delaying the input image signal by one frame, and a comparator 302.

An image signal supplied from the image input unit 101 in Fig. 11 is applied to the comparator 302 directly via a path P(x,y) and indirectly via a path Q(x,y) through the frame delay circuit 301. The latter is a one-pixel delay circuit corresponding to one frame of the image signal. Accordingly, sets of pixels are supplied to the comparator 302 sequentially, each set comprising two pixels at identical positions in the immediately preceding frame and present frame. The comparator 302 compares the value of the pixel of the preceding frame with the value of the pixel of the present frame, determines whether the difference between the compared values exceeds a predetermined quantity and outputs the detection signal 110 if the predetermined quantity is exceeded. More specifically, the comparator 302 detects motion between frames in the progressive image signal and applies the result of detection to the region designation unit 106.

According to the fifth embodiment, therefore, the detection signal 110 is output as a high-level signal if the following relation holds:

$$\text{abs}\{Q(x,y) - P(x,y)\} > K \quad \dots (14)$$

where K represents a predetermined value. It should be noted that $\text{abs}\{(x,y) - P(x,y)\}$ in Equation (14) indicates the absolute value of the difference between the values of pixel Q(x,y) and pixel P(x,y).

In accordance with the fifth embodiment as described above, motion of an image can be detected automatically based upon the difference between pixel values from one frame of an image to the next, thereby making it possible to select an image region to undergo highly efficient encoding.

[Sixth Embodiment]

A block-based motion detection method is well known from the MPEG standard, etc., as a motion detection method other than those described above. The construction of an encoding apparatus using a motion detector that employs this block-based motion detection method also is covered by the scope of the present invention.

Fig. 12 is a block diagram illustrating an example of the construction of a motion vector detector according to a sixth embodiment.

As shown in Fig. 12, the motion vector detector 107b includes a block forming unit 901, a motion vector calculation unit 902 and a comparator 903. The image signal supplied from the image input unit 101 is split into blocks each comprising 8×8 pixels by the block forming unit 901. The motion vector calculation unit 902 calculates a vector (u,v) , which indicates, with regard to each individual block of the blocks output from the block forming unit 901, the position of the

block relative to another block that has the highest degree of correlation. The comparator 903 compares the magnitude $[\sqrt{(u^2+v^2)}]$ of the vector (u,v) supplied from the motion vector calculation unit 902 with a first
5 predetermined value **a** and a second predetermined value **b**, determines that significant motion regarding this block has been verified if the magnitude $[\sqrt{(u^2+v^2)}]$ of the vector is greater than the first predetermined value **a** and is equal to or less than the second predetermined
10 value **b**, and outputs the detection signal 110 as a high level. More specifically, with regard to each block of pixels, the comparator 903 detects suitable motion defined by the predetermined upper and lower limit values **a**, **b** and supplies the result of detection to the
15 region designation unit 106.

Thus, the region designation unit 106 receives the detection signal 110 from the motion detector 107 (107a, 107b) and, when the target image is subjected to the discrete wavelet transform, generates the region
20 information 111 indicating which coefficients belong to the region in which motion has been detected and supplies the region information 111 to the quantizer 103.

The quantizer 103 quantizes the sequence of
25 coefficients supplied from the discrete wavelet transformation unit 102. At this time the region that

has been designated by the region information 111 from
the region designation unit 106 is quantized upon
shifting up the output of the quantizer 103 a
predetermined number of bits or raising quantization
5 precision a predetermined amount, this region of the
image is compared with the image periphery and is
encoded to a higher image quality. The output of the
quantizer 103 thus obtained is supplied to the entropy
encoder 104.

10 The entropy encoder 104 decomposes the data
sequence from the quantizer 103 into bit planes, applies
binary arithmetic encoding on a per-bit-plane basis and
supplies the code output unit 105 with a code sequence
indicative of the result of encoding. It should be
15 noted that a multilevel arithmetic encoder that does not
decompose data into bit planes or a Huffman encoder may
be used to construct the entropy encoder without
detracting from the effects of the present invention.
Such an encoder also is covered by the scope of the
20 present invention.

By virtue of this arrangement, a region of motion
within an image is encoded to an image quality higher
than that of the image periphery. This is to deal with
video shot by a surveillance video camera or by a
25 substantially stationary video camera that shoots
everyday scenes. In most cases the main item of

interest in such captured video resides in the region of the image where there is motion. By adopting the above-described arrangement, therefore, the portion of the image where the main item of interest appears can be
5 encoded to an image quality that is higher than that of the other regions of the image such as the background thereof.

The converse arrangement, namely one in which an image region in which motion is not detected is
10 designated as a target region for encoding at a higher performance, also is considered to be included as an embodiment of the present invention. Such an arrangement may be so adapted that a region for which the detection signal 110 is at the low level in the each
15 of the foregoing embodiments is made the object of highly efficient encoding. With such an arrangement, a region exhibiting little motion in an image will be encoded more efficiently than other regions.

For example, consider video shot by a video camera
20 tracking a moving subject such as an athlete. Here the background is detected as moving while the athlete being tracked by the camera exhibits little motion. By designating the image region in which motion is not detected as a region to undergo highly efficient
25 encoding, an athlete that is the subject of photography

in a sports scene can be encoded to an image quality higher than that of the background.

In accordance with the sixth embodiment, as described above, a region exhibiting motion in an image
5 can be detected automatically and an image region to which highly efficient encoding is to be applied can be selected.

[Seventh Embodiment]

The present invention covers also an arrangement in
10 which whether the region designation unit 106 outputs the region information 111 for a region in which motion has been detected or for a region in which motion has not been detected is switched in conformity with a change in the shooting conditions.

15 Fig. 13 is a block diagram illustrating an example of the construction of a region designation unit 106a according to the seventh embodiment for changing over the region designating operation automatically in dependence upon the shooting conditions.

20 As shown in Fig. 13, the region designation unit 106a includes a counter 904, a comparator 905, an inverter 906 and a changeover unit 907.

The detection signal 110 indicative of the result of detection by the motion detector 107 (107a, 107b) is
25 applied to the counter 904. On the basis of the detection signal 110, the counter 904 counts the number

of pixels contained in the image region that exhibits the detected motion. The detection signal 110 varies pixel by pixel, as described earlier. Therefore, by counting the number of times the detection signal 110 changes in accordance with the level of the detection signal 110, the number of pixels contained in the image region that exhibits motion can be obtained. The comparator 905 compares the pixel count obtained by the counter 904 with a predetermined value and applies a control signal 910 to the changeover unit 907. The changeover unit 907 is further supplied with the detection signal 110 indicative of the region in which motion has been detected, and a signal obtained by inverting the detection signal 110 by the inverter 906, namely a signal indicative of an image region in which motion has not been detected. If the comparator 905 finds that the number of pixels in the image region in which motion has been detected is equal to or less than a predetermined value, the changeover unit 907 selects and outputs the region information 111 based upon the detection signal 110 indicative of the image region in which motion has been detected. Conversely, if the comparator 905 finds that the number of pixels counted is greater than the predetermined value, then the changeover unit 907 selects and outputs the region information 111 based upon the output of the inverter

906, namely the image region in which motion has not been detected, in order that this region will be encoded to a high image quality.

In accordance with the seventh embodiment, as
5 described above, a region in which image motion has been detected or a region in which image motion has not been detected can be selected in dependence upon the characteristics of the image as an image region to undergo highly efficient encoding.

10 The encoding apparatus according to the above-described embodiment is effective for highly efficient encoding of a moving image. However, by treating a single image in a sequence of moving images as a still image, the apparatus can be applied to highly efficient
15 encoding of this still image. Such an arrangement also is covered by the scope of the present invention.

In each of the foregoing embodiments, hardware implementation of the various components is taken as an example. However, this does not impose a limitation
20 upon the present invention because the operations of these components can be implemented by a program executed by a CPU.

Further, though each of the foregoing embodiments has been described independently of the others, this
25 does not impose a limitation upon the invention because

the invention is applicable also to cases where these embodiments are suitably combined.

The present invention can be applied to a system constituted by a plurality of devices (e.g., a host
5 computer, interface, reader, printer, etc.) or to an apparatus comprising a single device (e.g., a copier or facsimile machine, etc.).

Furthermore, it goes without saying that the object of the invention is attained also by supplying a storage
10 medium (or recording medium) storing the program codes of the software for performing the functions of the foregoing embodiments to a system or an apparatus, reading the program codes with a computer (e.g., a CPU or MPU) of the system or apparatus from the storage
15 medium, and then executing the program codes. In this case, the program codes read from the storage medium implement the novel functions of the embodiments and the storage medium storing the program codes constitutes the invention. Furthermore, besides the case where the
20 aforesaid functions according to the embodiments are implemented by executing the program codes read by a computer, it goes without saying that the present invention covers a case where an operating system or the like running on the computer performs a part of or the
25 entire process in accordance with the designation of program codes and implements the functions according to

the embodiment.

It goes without saying that the present invention further covers a case where, after the program codes read from the storage medium are written in a function expansion card inserted into the computer or in a memory provided in a function expansion unit connected to the computer, a CPU or the like contained in the function expansion card or function expansion unit performs a part of or the entire process in accordance with the designation of program codes and implements the function of the above embodiments.

Thus, in accordance with the embodiments as described above, a region of interest can be extracted from multilevel image data at high speed. As a result, it is possible to provide an image encoding method and apparatus capable of adaptively selecting encoding processing that differs for each region of an image.

The present invention is not limited to the above embodiments and various changes and modifications can be made within the spirit and scope of the present invention. Therefore, to apprise the public of the scope of the present invention, the following claims are made.